

A General Method for Geometric Feature Matching and Model Extraction

Clark F. Olson

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Mail Stop 125-209, Pasadena, CA 91109-8099
<http://robotics.jpl.nasa.gov/people/olson/homepage.html>

January 4, 2000

Abstract

Popular algorithms for feature matching and model extraction fall into two broad categories, generate-and-test and Hough transform variations. However, both methods suffer from problems in practical implementations. Generate-and-test methods are sensitive to noise in the data. They often fail when the generated model fit is poor due to error in the selected features. Hough transform variations are somewhat less sensitive to the noise, but implementations for complex problems suffer from large time and space requirements and the detection of false positives. This paper describes a general method for solving problems where a model is extracted from or fit to data that draws benefits from both generate-and-test methods and those based on the Hough transform, yielding a method superior to both. An important component of the method is the subdivision of the problem into many subproblems. This allows efficient generate-and-test techniques to be used, including the use of randomization to limit the number of subproblems that must be examined. Each subproblem is solved using pose space analysis techniques similar to the Hough transform, which lowers the sensitivity of the method to noise. This strategy is easy to implement and results in practical algorithms that are efficient and robust. We apply this method to object recognition, geometric primitive extraction, robust regression, and motion segmentation.

1 Introduction

The generate-and-test paradigm is a popular strategy for solving model matching problems such as recognition, detection, and fitting. The basic idea of this method is to generate (or predict) many hypothetical model positions using the minimal amount of information necessary to identify unique solutions. A sequence of such positions is tested, and the positions that meet some criterion are retained. Examples of this technique include RANSAC [19] and the alignment method [26].

The primary drawback to generate-and-test paradigm is sensitivity to noise. Let us call the features that are used in predicting the model position for some test the *distinguished features*, since they play a more important role in whether the test is successful. The other features are *undistinguished features*. Error in the distinguished features causes the predicted position to be in error. As the error grows, the testing step becomes more likely to fail.

To deal with this problem, methods have been developed to propagate errors in the locations of the distinguished features [2, 21]. Under the assumption of a bounded error region for each of the distinguished image features, these methods can place bounds on the locations to which the undistinguished model features can be located in an image. When we count the number of undistinguished model features that can be aligned with image features (with the constraint that the distinguished features must always be in alignment up to the error bounds) these techniques can guarantee that we never undercount the number of alignable features. The techniques will thus never report that the model is not present according to some counting criterion when, in fact, the model does meet the criterion.

On the other hand, this method is likely to overcount the number of alignable features, even if the bounds on the location of each individual feature are tight. The reason for this is that, while this method checks whether there is a model position that brings each of the undistinguished model features into alignment with image features (along with all of the distinguished features) up to the error bounds, it does not check whether there is a position that brings all of the counted undistinguished features into alignment up to the error bounds.

A competing technique for feature matching and model extraction is based on the Hough transform. This method also generates hypothetical model positions solutions using minimal information, but rather than testing each solution separately, the testing is performed by analyzing the locations of the solutions in the space of possible model positions (or *poses*). This is often, but not always, accomplished through a histogramming or clustering procedure. The large clusters in the pose space indicate good model fits. We call techniques that examine the pose space for sets of consistent matches among all hypothetical matches *Hough-based methods*, since they derive from the Hough transform [28, 36]. While these techniques are less sensitive to noise in the features, they are prone to large computational and memory requirements, as well as the detection of false positive instances [20], if the pose space analysis is not careful.

In this paper, we describe a technique that combines the generate-and-test and Hough-based methods in a way that draws ideas and advantages from each, yielding a method that improves upon both. Like the generate-and-test method, (partial) solutions based on distinguished features are generated for further examination. However, each such solution is under-constrained and Hough-based methods are used to determine and evaluate the remainder of the solution. This allows both randomization to be used to reduce the computational complexity of the method and error propagation techniques to be used in order to better extract the relevant models. We call this technique RUDR (pronounced “rudder”), for Recognition Using Decomposition and Randomization.

First, it is shown that the problem can be treated as many subproblems, each of which is much simpler than the original problem. We next discuss various methods by which the

subproblems can be solved. The application of randomization to reduce the number of subproblems that must be examined is then described. These techniques yield efficiency gains over conventional generate-and-test and Hough-based methods. In addition, the subdivision of the problem allows us to examine a much smaller parameter space in each of the subproblems than in the original problem and this allows the error inherent in localization procedures to be propagated accurately and efficiently in the matching process.

This method has a large number of applications. It can be applied to essentially any problem where a model is fit to cluttered data (i.e. with outliers or multiple models present). We discuss the application of this method to object recognition, curve detection, robust regression, and motion segmentation.

The work described here is a generalization of previous work on feature matching and model extraction [44, 45, 47]. Similar ideas have been used by other researchers. A simple variation of this method has been applied to curve detection by Murakami et al. [43] and Leavers [35]. In both of these cases, the problem decomposition was achieved through the use of a single distinguished feature in the image for each of the subproblems. We argue that the optimal performance is achieved when the number of distinguished features is one less than the number necessary to fully define the model position in the errorless case. This has two beneficial effects. First, it reduces the amount of the pose space that must be considered in each problem (and the combinatorial explosion in the sets of undistinguished features that are examined). Second, it allows a more effective use of randomization in reducing the computational complexity of the method. A closely related decomposition and randomization method has been described by Cass [12] in the context of pose equivalence analysis. He uses a *base match* to develop an approximation algorithm for feature matching under uncertainty.

2 Related research

In this section, we review previous research on generate-and-test and Hough-based algorithms for model matching.

2.1 Generate-and-test methods

The basic idea in generate-and-test methods is to iteratively generate hypothetical model positions in the data and test the positions. The hypothetical model positions are often generated by conjecturing matches between the model and a set of features in the data and then determining the model positions such that the model is aligned with the data features. These positions are then tested by comparing the position of the model against the remainder of the data. One common measure of the quality of the position is the number of data features that agree well with the model at this position.

The first use of the generate-and-test strategy for recognition in computer vision was in Roberts' seminal thesis [52]. Roberts used sets of point matches between an object and an image to determine the best least-squares fit of the model to the image under rigid motion

and perspective projection. If the fit was good enough then the rest of the model was transformed and compared to the image in an attempt to ensure that the match was correct. Heuristic methods were used to determine which point sets to compare.

Fischler and Bolles described the RANSAC (for RANdom SAMple Consensus) technique [19]. They suggested that, when solving for the model position, it is best to use the minimum number of data features necessary to yield a finite set of solutions (assuming no error). This reduces the likelihood that one of the data features does not belong to the model, since there may be outliers or multiple models present. They also used a randomization technique, where sets of data features are sampled randomly until the probability of at least one sample being correct is sufficiently large, assuming that the model is actually present in the image. The application of these techniques to pose determination from point features is the primary subject of [19]. These techniques have also been applied to curve and surface detection [8].

An alternative to random sampling of the sets of data features is the use of features that are distinctive in some manner. Bolles and Cain [7] capitalize on such distinctive features in their *local-feature-focus* method. Lowe [41] finds sets of features that are distinctive based on their proximity, parallelism, or collinearity (implying that they are more likely to belong to a single model). There have been many similar techniques for *grouping* feature sets that are similarly distinctive (a review can be found in [57], also see [29]).

A technique that can be used to speed up generate-and-test methods, in cases where the model consists of a set of discrete features, is feature set indexing (e.g. [13, 33, 34, 42]). These methods preprocess the set of model features such that, at run-time, the sets of model features that could have yielded a set of image features can be determined quickly. The method of recognizing an object using an election [34] allows the testing step to be performed efficiently as follows. Given a set of image features that are conjectured to belong to an object (denote this the *basis set*), form all of the sets of image features that include the basis set and one additional feature and index the sets of model features that could feasibly match them. Then, count the number of votes for each possible basis set by from the indexed sets of model features. This simultaneously measures the quality of matching each possible basis set in the model to the given image basis set in terms of the number of additional image features that are aligned well by the match of the image basis set to the model basis set.

One problem that simple generate-and-test methods can have is that noisy data features can cause errors in the initial position estimate that result in poor model detection. An alternative to simply testing the model position corresponding to some hypothesized data set is to iteratively refine the model position as additional data features are conjectured to belong to the model [3, 17, 41]. This can significantly improve the overall match, particularly when the initial data set is noisy, although it will not help if the initial matching contains an outlier. Another technique that can be used if the localization error can be modeled is to carefully propagate the effects of localization error in the testing step [1, 21]. While this technique will not miss instances of a model that satisfy some error criterion, it will find instances that do not satisfy the criterion.

In addition to the applications described above, generate-and-test techniques have been

applied to many other applications, such as least median of squares regression [55], object recognition [26], extraction of geometric primitives such as curves and surfaces [54], and motion segmentation [66].

2.2 Hough-based methods

Parameter space analysis techniques can be traced back to the patent of the Hough transform [22]. The Hough transform was initially used to track particle curves in bubble-chamber imagery. Subsequent work by Rosenfeld [53] and Duda and Hart [15] played a substantial role in popularizing the Hough transform in the image processing and computer vision community and it has since become an established technique for the detection of curves and surfaces, as well as many other applications. Surveys of Hough transform techniques and applications can be found in [28, 36]. The basic idea is that each data feature can be mapped into a manifold in the parameter space of possible curves (or model positions, in general). Typical implementations consider a quantized parameter space and count the number of data features that map to a manifold intersecting each cell in the quantized space. Cells with high counts correspond to curves present in the image.

Ballard demonstrated that the Hough transform can be generalized to detect arbitrary shapes (composed of discrete points) in images [4]. This method determines a mapping between image features and an accumulator for a parameter space describing the possible positions of the object. When each image feature has been mapped into the parameter space, instances of the shape yield local maxima in the accumulator. Ballard used pixel orientation information to speed up the algorithm and improve accuracy. Ballard suggested that pairs of features could be mapped into the parameter space to reduce the effort in mapping features into the parameter space, but considered this infeasible for most cases.

Relatively recent work has capitalized on the idea of mapping multiple features into the parameter space [5, 10, 35, 68]. The primary benefit that is gained is that only those features sets that are large enough to map to a single point (or a finite set of points) in the parameter space are examined. A drawback is that there are many such sets of features. An additional technique that has proven useful for improving the efficiency of the Hough transform is randomization [5, 32, 35, 68].

Generalized Hough transform techniques have been applied to many object recognition problems (and these techniques are sometimes called pose clustering in this context) [14, 38, 60, 61, 62, 63, 64, 67]. Most of these applications consider feature sets that map to single points in the parameter space of possible model transformations and then perform some clustering method to detect objects, often multi-dimensional histogramming. A number of alternative methods for analyzing the parameter space have been proposed. See, for example, [27, 37, 49] for techniques to improve the standard Hough transform. Some interesting methods that have been applied to object recognition can be found in [9, 11, 56].

Hough-based techniques have been applied to many applications, such as curve matching [31, 67], segmentation of moving objects [6, 18, 30, 51, 65], data compression [59], and determination of viewing distance and angle of gaze from stereo disparity [50].

3 General problem formalization

The class of problems that we attack using RUDR are those that require a model to be fit to a set of observed data features, where a significant portion of the observed data may be outliers (in fact, the model may contain outliers as well) or there may be multiple models present in the data. These problems can, in general, be formalized as follows.

Given:

- \mathcal{D} : The data to match. This data consists of a set of features or measurements, $\{\delta_1, \dots, \delta_d\}$, that have been extracted, for example, from an image. For simplicity, we assume that all of the data features are of a single geometric class, such as points or segments, but this restriction can be removed.
- \mathcal{M} : The model to be fit. This model may be a set of distinct features as is typical in object recognition, or it may be a parameterized manifold such as a curve or surface, as in geometric primitive extraction and robust regression. The only constraint on the model is that we must be able to determine hypothetical model positions by matching the model with data features.
- \mathcal{T} : The possible positions or transformations of the model. We use τ to denote individual transformations in this space.
- $A(\mathcal{M}, \mathcal{D}, \mathcal{T}, \tau, D)$: A binary-valued acceptance criterion that specifies whether a transformation, τ , satisfactorily brings the model into agreement with a set of data features, $D \in \mathcal{D}$. We allow this criterion to be a function of the full set of data features and the set of transformations to allow the criterion to select the single best subset of data features according to some criterion or to take into account global matching information.

Determine and report:

- All maximal sets of data features, $D \in \mathcal{D}$, for which there is a transformation, $\tau \in \mathcal{T}$, such that the acceptance criterion, $A(\mathcal{M}, \mathcal{D}, \mathcal{T}, \tau, D)$, is satisfied. (Only the maximal sets are reported so that the subsets of each maximal set need not be reported.)

This formalization is very general. Many problems can be formalized in this manner, including object recognition, geometric primitive extraction, motion segmentation, and robust regression.

A useful acceptance criterion is based on bounding the fitting error between the model and the data. Let $C(\mathcal{M}, \delta, \tau)$ be a function that determines whether the specified position of the model fits the data feature δ (e.g. up to a bounded error). We let $C(\mathcal{M}, \delta, \tau) = 1$, if the criterion is satisfied, and $C(\mathcal{M}, \delta, \tau) = 0$, otherwise. The model is said to be brought into alignment with a set of data features, $D = \{\delta_1, \dots, \delta_x\}$ up to the error criterion, if all of the individual features are brought into alignment:

$$\prod_{i=1}^x C(\mathcal{M}, \delta_i, \tau) = 1 \quad (1)$$

The bounded-error acceptance criterion specifies that a set of data features, $D = \{\delta_1, \dots, \delta_x\}$, should be reported, if the cardinality of the set meets some threshold ($x \geq c$), there is a

position of the model that satisfies (1), and the set is not a subset of some larger set that is reported.

While this criterion cannot incorporate global information, such as mean-square-error or least-median-of-squares, RUDR is not restricted to using this bounded-error criterion. This method has been applied to least-median-of-squares regression with excellent results [44].

Example 1 As a running example, we will consider the detection of circles in two-dimensional image data. For this case, our model, \mathcal{M} , is simply the parameterization of a circle, $(x - x_c)^2 + (y - y_c)^2 = r^2$, and our data, \mathcal{D} , is a set of image points. The space of possible transformations is the space of circles, $\mathcal{T} = [x_c, y_c, r]^T$. We use a bounded-error acceptance criterion such that a point is considered to be on the circle if $|\sqrt{(x - x_c)^2 + (y - y_c)^2} - r| < \epsilon$. We will report the circles that have $\sum_{i=1}^d C_\epsilon(\mathcal{M}, \delta_i, \tau) > \pi r$. In other words, we search for the circles that have half of their perimeter present in the image.

Example 2 For a second example, we will discuss object recognition under similarity transformations. Here, our model consists of a set of model points in two dimensions and our data consists of similar points extracted from an image. We consider all similarity transforms of the model (translation, rotation, and scale). Our acceptance criterion will be based on the fraction of the model points that are aligned with extracted image points up to the bounded error ϵ , so we will report the instances where $\sum_{i=1}^m C_\epsilon(\mathcal{M}, \delta_i, \tau) > fm$. (This makes an assumption that correspondence between image points and model points is one-to-one, but it is simple to remove this assumption in practice [24].)

4 Approach

Let us call the hypothetical correspondence between a set of data features and the model a *matching*. The generate-and-test paradigm and many Hough-based strategies solve for hypothetical model positions using matchings of the minimum cardinality to constrain the model position up to a finite ambiguity (assuming errorless features). We call the matchings that contain this minimal amount of information the *minimal matchings* and we denote their cardinality k . We consider two types of models. One type of model consists of a set of discrete features similar to the data features. The other is a parameterized model such as a curve or surface. When the model is a set of discrete features, the minimal matchings specify the model features that match each of the data features in the minimal matching and we call these *explicit matchings*. Otherwise, the data features are matched implicitly to the parameterized model and we thus call these *implicit matchings*.

In the generate-and-test paradigm, the model positions generated using the minimal matchings are tested by determining how well the undistinguished features are fit according to the predicted model position. In Hough-based methods, it is typical to determine the positions of the model that align each of the minimal matchings and detect clusters of these

positions in the parameter space that describes the set of possible model positions¹, but other pose space analysis techniques can be used (e.g. [9, 12]).

The approach that we take draws upon both generate-and-test techniques and Hough-based techniques. The underlying matching method may be any one of several pose space analysis techniques in the Hough-based method (see Section 4), but unlike previous Hough-based methods, the problem is subdivided into many smaller problems, in which only a subset of the minimal matchings is examined. When randomization is applied to selecting which subproblems to solve, a low computational complexity can be achieved with a low probability of failure.

The key to this method is to subdivide the problem into many small subproblems, in which a *distinguished matching* of some cardinality $g < k$ between data features and the model is considered. Only those minimal matchings that contain the distinguished matching are examined in each subproblem and this constrains the portion of the pose space that the subproblem considers. We could consider each possible distinguished matching of the appropriate cardinality as a subproblem, but we shall see that this is not necessary in practice.

4.1 Problem equivalence

Let's consider the effect of this decomposition of the problem on the matchings that are detected by a system using a bounded-error criterion, $C(\mathcal{M}, d, t)$, as described above. For now, we assume that we have some method of determining precisely those sets of data features that should be reported according to the bounded-error acceptance criterion. The implications of performing matching only approximately and the use of an acceptance criterion other than the bounded-error criterion are discussed subsequently.

Proposition 1: For any transformation, $\tau \in \mathcal{T}$, the following statements are equivalent:

1. Transformation τ brings at least x data features into alignment with the model up to the error criterion.
2. Transformation τ brings at least $\binom{x}{k}$ sets of data features with cardinality k into alignment with the model up to the error criterion.
3. For any distinguished matching of cardinality g that is brought into alignment with the model up to the error criterion by τ , there are $\binom{x-g}{k-g}$ minimal matchings that contain the distinguished matching that are brought into alignment up to the error criterion by τ .

Proof : The proof follows from combinatorics. We sketch the proof that (a) Statement 1 implies Statement 2, (b) Statement 2 implies Statement 3, and (c) Statement 3 implies Statement 1. The statements are thus equivalent.

(a) From Statement 1, there are at least x data features with $C(\mathcal{M}, \delta_i, \tau) = 1$. We can thus form at least $\binom{x}{k}$ distinct sets of these data features with cardinality k . Each such set has $\prod_{i=1}^k C(\mathcal{M}, \delta_i, \tau) = 1$. These matchings thus contribute at least $\binom{x}{k}$ to the sum.

¹Early Hough transform strategies mapped single features into manifolds in the parameter space, but further work in Hough transforms has improved on these techniques by mapping sets of data features into points in the parameter space [5, 10, 35, 68].

(b) To form the $\binom{x}{k}$ sets of data features that are brought into alignment with the model, we must have x individual data features satisfying $C(\mathcal{M}, \delta_i, \tau) = 1$. (If there were $y < x$ such features then we could only form $\binom{y}{k}$ minimal sufficient matchings satisfying Equation (1).) Choose any subset, \mathcal{G} , of these matches of cardinality g . Form the $\binom{x-g}{k-g}$ subsets of cardinality $k - g$ that do not include any feature in \mathcal{G} . Each of these subsets when combined with \mathcal{G} forms a minimal sufficient matching that is brought into alignment up to the error criterion since each of the individual features satisfies $C(\mathcal{M}, \delta_i, \tau) = 1$.

(c) From Statement 3, the g data features in the distinguished matching are brought into alignment up to the error criterion by τ . In addition there must exist $x - g$ additional data features that are brought into alignment up to the error criterion by τ to form the $\binom{x-g}{k-g}$ subsets of cardinality $k - g$ that are brought into alignment up to the error criterion by τ . Thus, in total, there must be $g + x - g = x$ data features that are brought into alignment up to the error criterion by τ .

This result indicates that as long as we examine one distinguished matching that belongs to each of the matchings that should be reported, the strategy of subdividing the problem into subproblems yields equivalent results to examining the original problem as long as the threshold on the number of matches is set appropriately.

This decomposition of the problem allows our method to be viewed as a class of generate-and-test methods, where distinguished matchings (rather than minimal matchings) are generated and the testing step is performed using a pose space analysis method (such as clustering or pose space equivalence analysis) rather than comparing a particular model position against the data.

4.2 Optimal cardinality

While distinguished matchings of any cardinality could be considered, we must balance the complexity of the subproblems with the number of subproblems that are examined. Increasing the cardinality of the distinguished matching is beneficial up to a point. As the size of the distinguished matching is increased, the number of minimal matchings that is examined in each subproblem is decreased and we have more constraint on the position of the model. The subproblems are thus simpler to solve.

By itself, this does not improve matters, since there are more subproblems to examine. However, since we use randomization to limit the number of subproblems that are examined, we can achieve a lower computational complexity by having more simple subproblems than fewer difficult ones. On the other hand, when we reach $g = k$, the method becomes equivalent to a generate-and-test technique and we lose both the benefits gained through the Hough-based analysis of the pose space and the property that the subproblems become simpler with larger distinguished matchings. We thus use distinguished matchings with cardinality $g = k - 1$.

4.3 Approximation algorithms

For practical reasons, we may not wish to use an algorithm that reports exactly those matchings that satisfy the error criterion, since such algorithms are often time consuming. In this case, we cannot guarantee that examining a distinguished matching that belongs to a solution that should be reported will result in detecting that solution. However, empirical evidence suggests that the examination of these subproblems yields superior results when an approximation algorithm is used [45], owing to failures that occur in the examination of full problem.

We can also use these techniques with acceptance criteria other than the bounded-error criterion. With other criteria, the proposition is no longer always true, but if an approximation algorithm is used to detect good matchings, examination of the subproblems often yields good results. For example, an application of these ideas to least-median-of-squares regression has yielded an approximation algorithm that is provably accurate with high probability, while previous approximation algorithms do not have this property [44].

Example 1 For our circle detection example, $k = 3$, since three points are sufficient to define a circle in the noiseless case. The above analysis implies that, rather than examining individual image features, or all triples of features, we should examine trials (or subproblems) where only the triples that share some distinguished pair of features in common. Multiple trials are examined to guard against missing a circle.

Example 2 For object recognition under similarity transformation, we have $k = 2$, since two feature matches are sufficient to determine the model position. In this case, we examine trials where the pairs of image points are examined that share a single distinguished point in common. For each such pair, we consider each pair of model points as a possible match.

5 Solving the subproblems

Now, we must use some method to solve each of the subproblems that are examined. We can use any method that determines the number of matchings of a given cardinality can be brought approximately into alignment with the model at a particular position.

5.1 Histogramming

The simplest method for solving the subproblems is to use multi-dimensional histogramming step in order to locate large clusters in the pose space. In this method, the parameter space is quantized and a counter is maintained for each cell in the quantized space. Each matching is mapped into a manifold in the parameter space and the counter associated with each parameter space cell that intersects that manifold is incremented. Clusters in the parameter space are found by locating cells with high counts. The primary reason this method is popular is that it requires linear time in the number of matchings that are considered, while

most other methods are more complex. While this method is fast if the parameter space is not too large and yields good results for many applications, it does not propagate the effects of error accurately, it introduces quantization effects, and it is time consuming for large parameter spaces.

To reduce the problem of a large parameter space, hierarchical decompositions of the parameter space have often been used [14, 27, 35, 37, 38, 62, 64]. This is usually performed through a coarse-to-fine search of the parameter space or by decomposing the parameter space along the orthogonal axes. Such decompositions of the parameter space allow the initial steps to be performed using only a few cells in the transformation space. Those that cannot lead to a large cluster are eliminated and the rest are examined in finer detail.

The problem of error propagation is more difficult to handle with this technique. For complex problems, it can become problematic to detect the clusters without also detecting a significant number of false positives [20]. One possibility is to discretize the pose space finely and to increment the counter for each cell that is consistent with each minimal matching (up to some error bounds) (see, for example, [58]). However, this requires much computation.

5.2 Pose constraint methods

Alternatively, recently developed pose equivalence analysis techniques developed by Breuel [9] and Cass [12] can be applied that allow localization error to be propagated accurately. The basic idea of these methods is that each match between a model feature and a data feature yields some constraints on the volume of the pose space that must be satisfied in order for the features to be aligned up to some error bound. These constraints divide the pose space into equivalence classes, within which the same sets of features are brought into correspondence. A search of the pose space is performed to locate the positions that satisfy the maximum number of constraints.

Cass [12] searches the pose space by examining the arrangement of the constraints using computational geometry techniques. Significant speedup is achieved through the use of an approximation algorithm. Breuel uses a method that adaptively divides the pose space, pruning cells that are not consistent with a sufficient number of the constraints. Unpruned cells are divided recursively until they are small enough to accept as valid model positions. Breuel’s experiments indicate that his techniques can operate in linear expected time in the number of matchings, so we can, in general, perform this step efficiently.

Both of these techniques can be used in conjunction with our approach to subdividing the problem. In the method of Cass, we would examine much less of the pose space in each trial, since only a few of the pose equivalence classes are consistent with each distinguished matching. In Breuel’s method, k -tuples are used to form closed constraint regions in the pose space, so our approach would examine only the k -tuples that contain the distinguished matching.

5.3 Subspace error propagation

In our method, only a small portion of the parameter space is examined in each subproblem. If it is assumed that there is no error in the data features in the distinguished matching, then each subproblem considers only a sub-manifold of the parameter space. In general, if there are p transformation parameters and each feature match yields b constraints on the transformation, then a subproblem where the distinguished matchings have cardinality g considers only a $(p - gb)$ -dimensional manifold of the transformation space in the errorless case². This allows us to parameterize the sub-manifold (using $p - gb$ parameters) and perform analysis in this lower dimensional space. A particularly useful case is when the resulting manifold has only one dimension (i.e. it is a curve). In this case, the subproblem can be solved very simply by parameterizing the curve and finding positions on the curve that are consistent with many minimal matchings.

When localization error in the data features is considered, the subproblems must (at least implicitly) consider a larger space than the manifold described above. The subproblems are still much easier to solve. A technique that is useful in this case is to project the set of transformations that are consistent with a minimal matching up to the error criterion onto the manifold that results in the errorless case and then perform clustering only in the parameterization of this manifold as discussed above [47].

This method slightly overestimates the total number of consistent matches, since matches that are consistent in the projections may not be in the full pose space. However, significant errors are unlikely, because the regions of the pose space consistent with each minimal matching do not deviate far from the manifold corresponding the errorless case, except in extreme circumstances (cases that are nearly degenerate).

Example 1 For circle detection, the circle positions that share a pair of points lie on a curve in the pose space. (The center of the circle is always on the perpendicular bisector of the two distinguished points.) We parameterize the positions using the signed distance d from the center of the circle to the midpoint between the distinguished points (positive if above, negative if below). This yields a unique descriptor for every circle containing the distinguished points. For each triple that is considered, we can project the pose space consistent with the triple onto the parameterization by considering which centers are possible given some error bounds on the point locations [47]. We determine if a circle is present in each trial by finely discretizing d and performing a simple Hough transform variation, where the counter for each bin is incremented for each triple that is consistent with the span represented by the counter. Figure 1 gives an example showing how the positions that are consistent with a minimal matching are found. Peaks in the accumulator are accepted if

²This is not always true. For example, consider the case where the data and the model consist of sets of three-dimensional points and the transformation space is the six-dimensional space of rigid motions. Each individual match between a data point and a model point yields three constraints on the position of the model. However, a pair of such matches yields only five constraints, since the rotation around the segment joining the points is unconstrained. In this case, the additional constraint lies in the distance between the points, which must be the same in both the model and the data.




Figure 1: We determine bounds on the position of the center of a circle passing through three points (up to the localization error) by examining the range of possible perpendicular bisectors for the segments between the points.

they surpass some predetermined threshold.

Example 2 In the recognition of point sets under similarity transformations, the positions consistent with a distinguished matching between a model point and an image point lie in a two-dimensional manifold in the four-dimensional pose space. If we translate both the model and the image to have their origin at the distinguished point, then the translation of the model is fixed by the distinguished matching (in the errorless case) and the search can be restricted to rotation and scale. In the subspace, we can determine the area consistent with each minimal matching containing the distinguished matching and use of a variation of the Hough transform to detect object locations. Alternatively, variations of the divide-and-conquer pose space techniques can be used [9, 25].

6 Randomization and complexity

A deterministic implementation of these ideas examines each possible distinguished matching with the appropriate cardinality. This requires $O(n^k)$ time, where n is the number of possible matches between a data feature and the model. When explicit matchings are considered, $n = md$, where m is the number of model features and d is the number of data features. When

implicit matchings are considered, $n = d$. Such a deterministic implementation performs much redundant work. There are many distinguished matchings that are part of each of the large consistent matchings that we are seeking. We thus find each matching that meets the acceptance criterion many times (once for each distinguished matching that is contained in the maximal matching). We can take advantage of this redundancy through the use of a common randomization technique to limit the number of subproblems that we must consider while maintaining a low probability of failure.

Assume that some minimum number of the image features belong to the model. Denote this number b . Since our usual acceptance criterion is based on counting the number of image features that belong to the model, we can allow the procedure to fail when too few image features belong to the model. Otherwise, the probability that some set of image features with cardinality $g = k - 1$ completely belongs to the model is approximately bounded by $\left(\frac{b}{d}\right)^{k-1}$. If we take t trials that select sets of $k - 1$ image features randomly, then the probability that none of them will completely belong to the model is:

$$p_t \approx \left(1 - \left(\frac{b}{d}\right)^{k-1}\right)^t. \quad (2)$$

Setting this probability below some arbitrarily small threshold ($p_t < \gamma$) yields:

$$t \approx \frac{\ln \gamma}{\ln(1 - \left(\frac{b}{d}\right)^{k-1})} \approx \left(\frac{d}{b}\right)^{k-1} \ln \frac{1}{\gamma}. \quad (3)$$

Now, for explicit matches, we assume that some minimum fraction f_e of the model features appear in the image. In this case, the number of trials necessary is approximately $\left(\frac{d}{f_e m}\right)^{k-1} \ln \frac{1}{\gamma}$. For each trial, we must consider matching the set of image features against each possibly matching set of model features, so the total number of distinguished matchings that are considered is approximately $\left(\frac{d}{f_e}\right)^{k-1} (k - 1)! \ln \frac{1}{\gamma}$. Each explicit distinguished matching requires $O(md)$ time to process, so the overall time required is $O(md^k)$.

For implicit matches, we may assume that each significant model in the image comprises some minimum fraction f_i of the image features. The number of trials necessary to achieve a probability of failure below γ is approximately $f_i^{1-k} \ln \frac{1}{\gamma}$, which is a constant independent of the number of model or image features. Since each trial can be solved in $O(d)$ time, the overall time required is $O(d)$.

Note that the complexity can be reduced further by performing subsampling among the minimal matchings considered in each trial. Indeed, $O(1)$ complexity is possible with some assumptions about the number of features present and the rate of errors allowable [5]. We have not found this further complexity reduction to be necessary in our experiments. However, it may be useful when the number of image features is very large.

Grouping techniques can also be used to improve the computational complexity of the algorithm. If we have some method that can determine feature sets that are more likely to belong to the same model, we can use these as a set of distinguished features in the image.

Such method can be used to reduce the likelihood of a false positive match in addition to reducing the computational complexity [46].

Example 1 Our circle detection case uses implicit matchings. If we assume that each circle that we wish to detect comprises at least $f_i = 5\%$ of the image data and require that the probability of failure is below $\gamma = 0.1\%$, then the number of trials necessary is 2764. Each trial considers the remaining $d - 2$ image features. Note that techniques considering all triples will surpass the number of triples considered here when $d > 53$.

Example 2 Object recognition uses explicit matchings. For this case, we may wish to locate each instance of an object (with probability 0.999) where at least 75% of the object appears in the image. If the model consists of 20 feature points and there are 60 features in the image, then we should examine 28 distinguished points in the image. For each such point, we consider each of the model features as a possible match and then perform recognition in the given subspace.

7 Comparison with previous techniques

This section gives a comparison of the RUDR approach with previous generate-and-test and Hough-based techniques.

Deterministic generate-and-test techniques require $O(n^{k+1})$ time to perform model extraction in general, since there are $O(n^k)$ minimal matchings and the testing stage can be implemented $O(n)$ time. This can often be reduced slightly through the use of efficient geometric searching techniques during the testing stage (e. g. [39]). RUDR yields a superior computational complexity requirement. When randomization is applied to generate-and-test techniques, the computation complexity becomes $O(md^{k+1})$ (or slightly better using efficient geometric search) for explicit matches and $O(d)$ for implicit matches. RUDR yields a superior computational complexity for the case of explicit matches and, while the generate-and-test approach matches the complexity for the case of implicit matches, RUDR examines less subproblems by a constant factor (approximately $\frac{1}{f_i}$) and is thus faster in practice.

In addition, previous generate-and-test techniques are inherently less precise in the propagation of localization error. The basic generate-and-test algorithm introduces false positives unless care is taken to propagate the errors correctly [2, 21], since error in the data features leads to error in the hypothetical model pose and this error causes some of the models to be missed as a result of a poor fit. On the other hand, when error propagation is used, false positives are introduced. The techniques ensure that each of the undistinguished features can be separately brought into alignment (along with the distinguished set) up to some error bounds by a single model position, this position may be different for each such feature match. It does not guarantee that *all* of the features can be brought into alignment up to the error bounds by a single position and thus causes false positives to be found.

Hough-based methods are capable of propagating localization error such that neither false positives nor false negatives occur (in the sense that only matchings meeting the acceptance criterion are reported) [9, 12]. However, previous Hough-based methods have had large time and space requirements. Deterministic Hough-based techniques that examine minimal matchings require $O(n^k)$ time and considerable memory [45].

Randomization has been previously applied to Hough transform techniques [5, 32, 35, 68]. However, in previous methods, randomization has been used in a different manner than it is used here. While our approach examines all of the data in each of the subproblems, previous uses of randomization in Hough-based methods subsample the overall data examined, causing both false positives and false negatives to occur as a result. While false negatives can occur due to the use of randomization in the our approach, the probability of such an occurrence can be set arbitrarily low.

Our method draws the ability to propagate localization error accurately from Hough-based methods and combines it with the ability to subdivide the problem into many smaller subproblems and thus reap the full benefit of randomization techniques. The result is a model extraction algorithm with superior computational complexity to previous methods that is also robust with respect to false positives and false negatives.

All of the techniques considered so far have been model-based methods. The primary drawback to such techniques is a combinatorial complexity that is polynomial in the number of features, but exponential in the complexity of the pose space (as measured by k). This can be subverted in some cases by assuming that some fraction of the data features arises from the model (this shifts the base of the exponent to the required fraction). An alternative that can be useful in reducing this problem is the use of grouping or perceptual organization methods that use data-driven techniques to determine features that are likely to belong to the same model (for example, [29, 40]). In cases where models can be identified by purely data-driven methods, such techniques are likely to be faster than the techniques described here. However, work has shown the even imperfect feature grouping methods can improve both the complexity and the rate of false positives in the RUDR method [46].

There are some situations where RUDR can not be applied effectively. If a single data feature is sufficient to constrain the position of the model, the RUDR problem decomposition will not be useful. In addition, the techniques we describe will be of less value is when there is a small number of features in the image. In this case, the randomization may not yield an improvement in the speed of the algorithm. However, the error propagation benefits will still apply.

8 Applications of RUDR

RUDR has been applied to several problems. We review the important aspects of these applications here and discuss additional areas where RUDR can be applied.

8.1 Extraction of geometric primitives

As previously discussed, the Hough transform is a well known technique for geometric primitive extraction [28, 36]. The application of RUDR to this method improves the efficiency of the technique, allows the localization error to be propagated accurately, and reduces the amount of memory that is required [47].

Consider the case of detecting curves from feature points in two-dimensional image data. If we wish to detect curves with p parameters, then we use distinguished matchings consisting of $p-1$ feature points, since, in general, p points are required to solve for the curve parameters. Each distinguished matching maps to a one-dimensional manifold (a curve) in the parameter space, if the points are errorless and in general position. Methods have been developed to map minimal matchings with bounded errors into segments of this curve for the case of lines and circles [47]. $O(d)$ time and space is required for curve detection with these techniques, where d is the number of data points extracted from the image.

Figure 2 shows the results of using RUDR to detect circles in a binary image of an engineering drawing. The results are very good, with the exception of circles found with a low threshold that are not perceptually salient. However, these circles meet the acceptance criterion specified, so this is not a failure of the algorithm.

The image in Figure 2 contains 9299 edge pixels. In order to detect circles comprising 4% of the image, RUDR examines 4318 trials and considers 4.01×10^7 triples. Contrast this to the 8.04×10^{11} possible triples. A generate-and-test technique using the same type of randomization examines 1.08×10^5 trials (1.00×10^9 triples) to achieve the same the same probability of examining a trial where the distinguished features belong to some circle, but will still miss circles due to the error in the features.

The robustness of this technique for line detection has been compared against other methods in a large number of synthetic images. Four methods were compared:

1. The RUDR paradigm with propagated localization error.
2. The RUDR paradigm without propagated localization error.
3. A method mapping pairs of points into the parameter space, but without decomposition into subproblems.
4. The standard Hough transform.

Figure 3 shows the results. For each method, the probability of detecting the single correct line segment present in the image is plotted versus the probability of finding a false positive (curved distractors were added to the images) for varying levels of the threshold used to determine which lines are detected.

The best performance is achieved by the RUDR paradigm with propagation of localization error into the parameter space. Interestingly, the RUDR paradigm fares poorly when localization error is not propagated carefully. It is thus crucial to propagate the localization into the parameter space when using the RUDR paradigm.

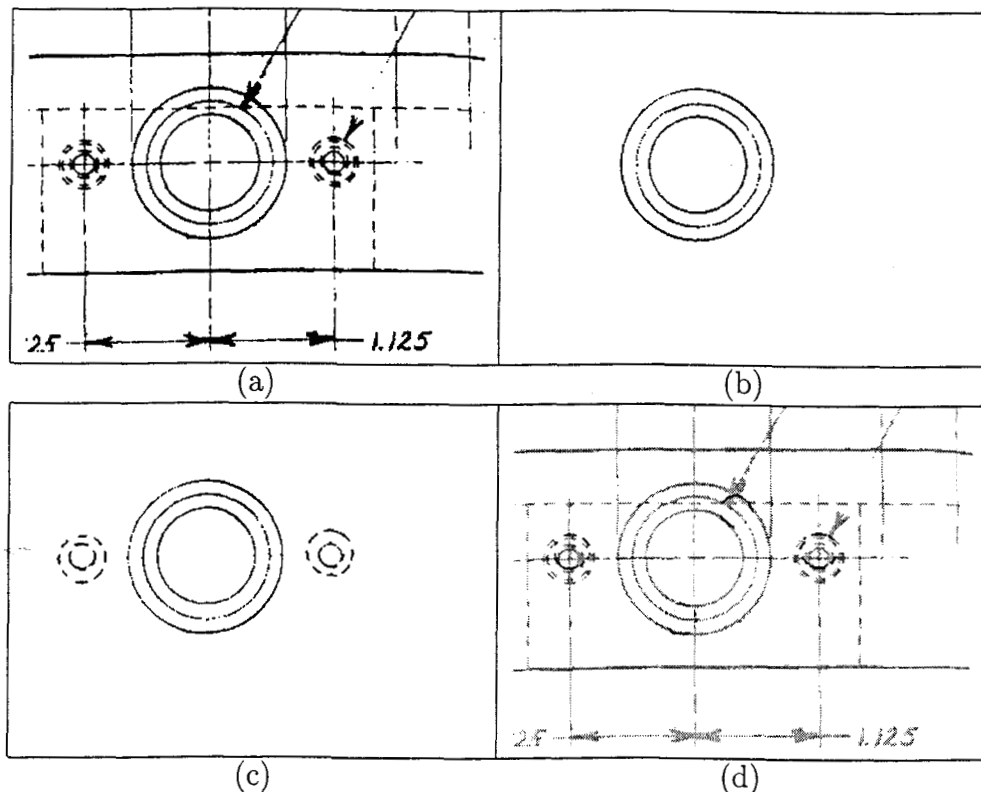


Figure 2: Circle detection. (a) Engineering drawing. (b) Circles found comprising 4% of the image. (c) Perceptually salient circles found comprising 0.8% of the image. (d) Insalient circles found comprising 0.8% of the image.

8.2 Robust regression

RUDR can be applied to the problem of finding the least-median-of-squares (LMS) regression line. We examine the problem of fitting a line to points in the plane. In this case, a single distinguished point is examined in each trial (since two points are required to define a line). For each trial, we determine the line that is optimal with respect to the median residual, but with the constraint that the line must pass through the distinguished point.

It can be shown that the solution to this constrained problem has a median residual that is no more than the sum of the optimal median residual and the distance of the distinguished point from the optimal LMS regression line [44]. Now, at least half of the data points must lie no farther from the optimal regression line than the optimal median residual (by definition). Each trial thus has a probability of at least 0.5 of obtaining a solution with a residual no worse than twice the optimal median residual. The use of randomization implies that we need to perform only a constant number of trials to achieve a good solution with high probability (approximately $-\log_2 \delta$ trials are necessary to achieve an error rate of δ).

Each subproblem (corresponding to a distinguished point) can be solved using a specialized method based on parametric search techniques [44]. This allows each subproblem




Figure 3: Receiver operating characteristic (ROC) curves for line detection generated using synthetic data.




Figure 4: Robust regression examples. The solid lines are the RUDR LMS estimate. The dashed lines are the PROGRESS LMS estimate. The dotted lines are the least-squares fit.

to be solved exactly in $O(n \log^2 n)$ time or in $O(n \log n)$ time for a fixed precision solution using numerical techniques. These techniques have also been extended to problems in higher dimensional spaces.

The complexity of our method is superior to the best known exact algorithms for this problem [16]. The PROGRESS algorithm [55] is a commonly used approximation algorithm for LMS regression that is based on the generate-and-test paradigm. It requires $O(n)$ time. However, unlike our algorithm, this algorithm yields no lower bounds (with high probability) on the quality of the solution detected.

Figure 4 shows two examples where RUDR, PROGRESS, and least-squares estimation were used to perform linear regression. In these examples, we used 400 inliers and 100 outliers, both from two-dimensional normal distributions. trial of the RUDR algorithm were considered, and 50 trials of the PROGRESS algorithm. For both examples, RUDR produces the best fit to the inliers. The least-squares fit is known to be non-robust, so it is not surprising that it fails poorly. The PROGRESS algorithm had difficulty, since, even in 50 trials, it failed generate a solution very close to the optimal solution.

8.3 Object recognition

The application of RUDR to object recognition yields an algorithm with $O(md^k)$ computational complexity, where m is the number of model features, d is the number of data features, and k is the minimal number of feature matches necessary to constrain the position of the model up to a finite ambiguity in the case of errorless features in general position.

We examine the recognition of three-dimensional objects using two-dimensional image data, for which $k = 3$. In each subproblem, we compute the pose for each minimal matching containing the distinguished matching using the method of Huttenlocher and Ullman [26]. We then use a multi-dimensional histogramming technique that examines each axis of the




Figure 5: Recognition of occluded two-dimensional objects. (a) The corners detected in the image. (b) The best hypotheses found for the occluded objects with the edges drawn in.

pose space separately. After finding the clusters along some axis in the pose space, the clusters of sufficient size are then analyzed recursively in the remainder of the pose space [45]. The poses for all sets of points sharing a distinguished matching of cardinality $k - 1$ lie in a two-dimensional subspace for this case. Despite this, we perform the histogramming in the full six-dimensional space, since this requires little extra time and space with this histogramming method. Feature error has been treated in an ad hoc manner in this implementation through the examination of overlapping bins in the pose space. Complex images may require a more thorough analysis of the errors.

We can also apply these techniques to images in which imperfect grouping techniques have determined sets of points that are likely to derive from the same object [46]. This allows a reduction in both the computational complexity and the rate of false positives.

Figure 5 shows an example where RUDR was used to recognize occluded two-dimensional figures by matching feature points. The partial occlusion of the objects and the distracting features from other objects did not prevent the method from recognizing the figures. Figure 6 shows an example where this approach has been applied to the recognition of a three-dimensional object.

8.4 Motion segmentation

RUDR can be used to perform motion segmentation with any technique for determining structure and motion from corresponding data features in multiple images. In this problem, we are given sets of data features in multiple images. We assume that we know the feature correspondences between images (e.g. from a tracking mechanism), but not which sets of features belong to coherent objects.




Figure 6: Three-dimensional object recognition. (a) Corners detected in the image. (b) Best hypothesis found.

Say that we have an algorithm to determine structure and motion using k feature correspondences in i images and that there are d features for which we know the correspondences between the images (see [23] for a review of such techniques). We examine distinguished matchings of $k - 1$ sets of feature correspondences between the images. Each subproblem is solved by determining the hypothetical structure and motion of each minimal matching (k sets of feature correspondences) containing the distinguished matching and then determining how many of the minimal matchings yield consistent structures for the distinguished matching and motions that are consistent with them belonging to a single object. This is repeated for enough distinguished matchings to find all of the rigidly moving objects consisting of some minimum fraction of all image features.

Our analysis for implicit matchings implies that we must examine approximately $\epsilon^{1-k} \ln \frac{1}{\gamma}$ trials to find objects whose fraction of the total number of data features is at least ϵ with a probability of failure for a particular object no larger than γ .

9 Summary

This paper has described a technique that we have named RUDR for solving model extraction and fitting problems such as recognition and regression. This approach is very general and can be applied to a wide variety problems where a model is fit to a set of data features and it is tolerant to noisy data features, occlusion, and outliers.

The RUDR method draws advantages from both the generate-and-test paradigm and from parameter space methods based on the Hough transform. The key ideas are:

1. Break down the problem into many small subproblems in which only the model positions consistent with some distinguished matching of features are examined.
2. Use randomization techniques to limit the number of subproblems that need to be examined to guarantee a low probability of failure.

3. Use clustering or parameter space analysis techniques to determine the matchings that satisfy the criteria.

The use of this technique yields two primary advantages over previous methods. First, RUDR is computationally efficient and has a low memory requirement. Second, we can use methods by which the localization error in the data features is propagated precisely, so that false positives and false negatives do not occur.

Acknowledgments

The research described in this paper was carried out in part by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored in part by the Air Force Research Laboratory at Tyndall Air Force Base, Panama City, Florida, through an agreement with the National Aeronautics and Space Administration. Portions of this research were carried out while the author was with Cornell University and the University of California at Berkeley. A condensed version of this work appeared in the Vision Algorithm '99 Workshop [48].

References

- [1] T. D. Alter and D. W. Jacobs. Error propagation in full 3d-from-2d object recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 892–898, 1994.
- [2] T. D. Alter and D. W. Jacobs. Uncertainty propagation in model-based recognition. *International Journal of Computer Vision*, 27(2):127–159, 1998.
- [3] N. Ayache and O. D. Faugeras. HYPER: A new approach for the recognition and positioning of two-dimensional objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):44–54, January 1986.
- [4] D. H. Ballard. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2):111–122, 1981.
- [5] J. R. Bergen and H. Shvaytser. A probabilistic algorithm for computing Hough transforms. *Journal of Algorithms*, 12:639–656, 1991.
- [6] M. Bober and J. Kittler. Estimation of complex multimodal motion: An approach based on robust statistics and Hough transform. *Image and Vision Computing*, 12(10):661–668, December 1994.
- [7] R. C. Bolles and R. A. Cain. Recognizing and locating partially visible objects: The local-feature-focus method. *International Journal of Robotics Research*, 1(3):57–82, 1982.

- [8] R. C. Bolles and M. S. Fischler. A RANSAC-based approach to model fitting and its application to finding cylinders in range data. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 637–642, 1981.
- [9] T. M. Breuel. Fast recognition using adaptive subdivisions of transformation space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 445–451, 1992.
- [10] A. Califano, R. M. Bolle, and R. W. Taylor. Generalized neighborhoods: A new approach to complex parameter feature extraction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 192–199, 1989.
- [11] T. A. Cass. Polynomial-time object recognition in the presence of clutter, occlusion, and uncertainty. In *Proceedings of the European Conference on Computer Vision*, pages 834–842, 1992.
- [12] T. A. Cass. Polynomial-time geometric matching for object recognition. *International Journal of Computer Vision*, 21(1/2):37–61, January 1997.
- [13] D. T. Clemens and D. W. Jacobs. Space and time bounds on indexing 3-d models from 2-d images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10):1007–1017, October 1991.
- [14] M. Dhome and T. Kasvand. Polyhedra recognition by hypothesis accumulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(3):429–438, May 1987.
- [15] R. O. Duda and P. E. Hart. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15:11–15, 1972.
- [16] H. Edelsbrunner and D. L. Souvaine. Computing least median of squares regression lines and guided topological sweep. *Journal of the American Statistical Association*, 85(409):115–119, March 1990.
- [17] O. D. Faugeras and M. Hebert. The representation, recognition, and locating of 3-d objects. *International Journal of Robotics Research*, 5(3):27–52, 1986.
- [18] C. L. Fennema and W. B. Thompson. Velocity determination in scenes containing several moving objects. *Computer Graphics, and Image Processing*, 9:301–315, 1979.
- [19] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24:381–396, June 1981.
- [20] W. E. L. Grimson and D. P. Huttenlocher. On the sensitivity of the Hough transform for object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(3):255–274, March 1990.

- [21] W. E. L. Grimson, D. P. Huttenlocher, and D. W. Jacobs. A study of affine matching with bounded sensor error. *International Journal of Computer Vision*, 13(1):7–32, 1994.
- [22] P. V. C. Hough. Method and means for recognizing complex patterns. U. S. Patent 3069654, 1962.
- [23] T. S. Huang and A. N. Netravali. Motion and structure from feature correspondences: A review. *Proceedings of the IEEE*, 82(2):252–268, February 1994.
- [24] D. P. Huttenlocher and T. A. Cass. Measuring the quality of hypotheses in model-based recognition. In *Proceedings of the European Conference on Computer Vision*, pages 773–775, 1992.
- [25] D. P. Huttenlocher and W. J. Rucklidge. A multi-resolution technique for comparing images using the Hausdorff distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 705–706, 1993.
- [26] D. P. Huttenlocher and S. Ullman. Recognizing solid objects by alignment with an image. *International Journal of Computer Vision*, 5(2):195–212, 1990.
- [27] J. Illingworth and J. Kittler. The adaptive Hough transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):690–698, September 1987.
- [28] J. Illingworth and J. Kittler. A survey of the Hough transform. *Computer Vision, Graphics, and Image Processing*, 44:87–116, 1988.
- [29] D. W. Jacobs. Robust and efficient detection of salient convex groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1):23–37, January 1996.
- [30] S. N. Jayaramamurthy and R. Jain. An approach to the segmentation of textured dynamic scenes. *Computer Vision, Graphics, and Image Processing*, 21:239–261, 1983.
- [31] A. Kalvin, E. Schonberg, J. T. Schwartz, and M. Sharir. Two-dimensional, model-based, boundary matching using footprints. *International Journal of Robotics Research*, 5(4):38–55, 1986.
- [32] N. Kiryati, Y. Eldar, and A. M. Bruckstein. A probabilistic Hough transform. *Pattern Recognition*, 24(4):303–316, 1991.
- [33] T. F. Knoll and R. C. Jain. Recognizing partially visible objects using feature indexed hypotheses. *IEEE Journal of Robotics and Automation*, 2(1):3–13, 1986.
- [34] Y. Lamdan, J. T. Schwartz, and H. J. Wolfson. Affine invariant model-based object recognition. *IEEE Transactions on Robotics and Automation*, 6(5):578–589, October 1990.

- [35] V. F. Leavers. The dynamic generalized Hough transform: Its relationship to the probabilistic Hough transforms and an application to the concurrent detection of circles and ellipses. *CVGIP: Image Understanding*, 56(3):381–398, November 1992.
- [36] V. F. Leavers. Which Hough transform? *CVGIP: Image Understanding*, 58(2):250–264, September 1993.
- [37] H. Li, M. A. Lavin, and R. J. Le Master. Fast Hough transform: A hierarchical approach. *Computer Vision, Graphics, and Image Processing*, 36:139–161, 1986.
- [38] S. Linnainmaa, D. Harwood, and L. S. Davis. Pose determination of a three-dimensional object using triangle pairs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):634–647, September 1988.
- [39] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM Journal on Computing*, 9(3):615–627, 1980.
- [40] D. G. Lowe. *Perceptual Organization and Visual Recognition*. Kluwer Academic, 1985.
- [41] D. G. Lowe. Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31:355–395, 1987.
- [42] R. Mehrotra and W. I. Grosky. Shape matching utilizing indexed hypotheses generation and testing. *IEEE Transactions on Robotics and Automation*, 5(1):70–77, 1989.
- [43] K. Murakami, H. Koshimizu, and K. Hasegawa. On the new Hough algorithms without two-dimensional array for parameter space to detect a set of straight lines. In *Proceedings of the IAPR International Conference on Pattern Recognition*, pages 831–833, 1986.
- [44] C. F. Olson. An approximation algorithm for least median of squares regression. *Information Processing Letters*, 63(5):237–241, September 1997.
- [45] C. F. Olson. Efficient pose clustering using a randomized algorithm. *International Journal of Computer Vision*, 23(2):131–147, June 1997.
- [46] C. F. Olson. Improving the generalized Hough transform through imperfect grouping. *Image and Vision Computing*, 16(9-10):627–634, July 1998.
- [47] C. F. Olson. Constrained Hough transforms for curve detection. *Computer Vision and Image Understanding*, 73(3):329–345, March 1999.
- [48] C. F. Olson. A general method for feature matching and model extraction. In *Proceedings of the Vision Algorithms Workshop*, 1999.
- [49] J. O’Rourke and K. R. Sloan. Dynamic quantization: Two adaptive data structures for multidimensional spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:266–279, 1984.

- [50] S. A. Peek, J. E. W. Mayhew, and J. P. Frisby. Obtaining viewing distance and angle of gaze from vertical disparity using a Hough-type accumulator. *Image and Vision Computing*, 2:180–190, 1984.
- [51] C. J. Radford. Optical flow fields in Hough transform space. *Pattern Recognition Letters*, 4:293–303, September 1986.
- [52] L. G. Roberts. *Machine Perception of Three-Dimensional Solids*. PhD thesis, Massachusetts Institute of Technology, May 1963.
- [53] A. Rosenfeld. *Picture Processing by Computer*. Academic Press, 1969.
- [54] G. Roth and M. D. Levine. Extracting geometric primitives. *CVGIP: Image Understanding*, 58(1):1–22, July 1993.
- [55] P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. John Wiley and Sons, 1987.
- [56] W. J. Rucklidge. Efficiently locating objects using the Hausdorff distance. *International Journal of Computer Vision*, 24(3):251–270, September 1997.
- [57] S. Sarkar and K. L. Boyer. Perceptual organization in computer vision: A review and a proposal for a classificatory structure. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2):382–399, March/April 1993.
- [58] S. D. Shapiro. Generalization of the Hough transform for curve detection in noisy digital images. In *Proceedings of the International Joint Conference on Pattern Recognition*, pages 710–714, 1978.
- [59] S. D. Shapiro. Use of the Hough transform for image data compression. *Pattern Recognition*, 12:333–337, 1980.
- [60] T. M. Silberberg, L. S. Davis, and D. A. Harwood. An iterative Hough procedure for three-dimensional object recognition. *Pattern Recognition*, 17(6):621–629, 1984.
- [61] T. M. Silberberg, D. A. Harwood, and L. S. Davis. Object recognition using oriented model points. *Computer Vision, Graphics, and Image Processing*, 35:47–71, 1986.
- [62] G. Stockman. Object recognition and localization via pose clustering. *Computer Vision, Graphics, and Image Processing*, 40:361–387, 1987.
- [63] G. Stockman, S. Kopstein, and S. Benett. Matching images to models for registration and object detection via clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(3):229–241, 1982.
- [64] D. W. Thompson and J. L. Mundy. Three-dimensional model matching from an unconstrained viewpoint. In *Proceedings of the IEEE Conference on Robotics and Automation*, volume 1, pages 208–220, 1987.

- [65] T. Y. Tian and M. Shah. Recovering 3d motion of multiple objects using adaptive Hough transform. In *Proceedings of the International Conference on Computer Vision*, pages 284–289, 1995.
- [66] P. H. S. Torr and D. W. Murray. Stochastic motion clustering. In *Proceedings of the European Conference on Computer Vision*, volume 2, pages 328–337, 1994.
- [67] J. L. Turney, T. N. Mudge, and R. A. Volz. Recognizing partially occluded parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(4):410–421, 1985.
- [68] L. Xu, E. Oja, and P. Kultanen. A new curve detection method: Randomized Hough transform (RHT). *Pattern Recognition Letters*, 11:331–338, May 1990.